

On Convolutional Approximations to Linear Dimensionality Reduction Operators for Large-Scale Data Processing

Swayambhoo Jain *Student Member, IEEE* and Jarvis Haupt *Member, IEEE*

Abstract

In this paper, we examine the problem of approximating a general linear dimensionality reduction (LDR) operator, represented as a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ with $m < n$, by a partial circulant matrix with rows related by circular shifts. Partial circulant matrices admit fast implementations via Fourier transform methods and subsampling operations; our investigation here is motivated by a desire to leverage these potential computational improvements in large-scale data processing tasks. We establish a fundamental result, that most large LDR matrices (whose row spaces are uniformly distributed) in fact cannot be well-approximated by partial circulant matrices. Then, we propose a natural generalization of the partial circulant approximation framework that entails approximating the range space of a given LDR operator $\mathbf{A} \in \mathbb{R}^{m \times n}$ over a restricted domain of inputs, using a matrix formed as a product of a partial circulant matrix having $m' > m$ rows and a $m \times m'$ “post-processing” matrix. We introduce a novel algorithmic technique, based on sparse matrix factorization, for identifying the factors comprising such approximations, and provide preliminary evidence to demonstrate the potential of this approach.

Index Terms

Subspace learning, circulant matrices, matrix factorization, sparse regularization, big data

I. INTRODUCTION

Numerous tasks in signal processing, statistics, and machine learning employ dimensionality reduction methods to facilitate the processing, visualization, and analysis of (ostensibly) high-dimensional data in (more tractable) low-dimensional spaces. Among the myriad of dimensionality reduction techniques in the literature, *linear dimensionality reduction* (LDR) methods remain among the most popular and widely-used.

One well-known example is principal component analysis [1] characterized by a linear dimensionality reduction (LDR) operator designed to maximally preserve the variance of the original data in the projected space, and which has been widely used for data compression, denoising, and as a dimensionality reducing pre-processing step for other

The authors are with the Department of Electrical and Computer Engineering at the University of Minnesota – Twin Cities. Email: {jainx174, jdaupt}@umn.edu. Manuscript submitted February 24, 2015.

This work was supported in part by DARPA/ONR Grant No. N66001-11-1-4090 and the DARPA Young Faculty Award, Grant No. N66001-14-1-4047.

analyses (e.g., clustering, classification, etc.) [2]. Other classical data analysis methods that employ specialized LDR operators include linear discriminant analysis (where the operator is designed to preserve separations among original data points belonging to disparate classes), and canonical correlations analysis (where the operator maximizes correlations among projected data points). See the survey paper [3] for additional examples.

In recent years, LDR methods have also been utilized for universal “precompression” in high-dimensional inference tasks. For example, fully random LDR operators are at the essence of the initial investigations into *compressed sensing* (CS) (see, e.g., [4]–[7]); and other, more structured, LDR operators – both non-adaptive (see, e.g., [8]–[17]) and adaptive (see, e.g., [18]–[32]) – have also been examined recently in the context of CS and sparse inference.

The computational efficiency of LDR methods is often cited as one of their primary virtues; however, as data of interest become increasingly large-scale (high-dimensional, and numerous), even the relatively low computational complexity associated with LDR methods can become significant. In this paper, we investigate the utility of employing *partial circulant approximations* to general LDR operators; partial circulant matrices admit fast implementations (via convolution or Fourier transform methods, and subsampling), and their use as surrogates to arbitrary LDR matrices may provide potentially significant computational efficiency improvements in practice.

A. Problem Statement and Our Contributions

We represent an arbitrary (real) linear dimensionality reduction (LDR) operator as a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, with $m < n$. Operating with \mathbf{A} on an arbitrary vector $\mathbf{x} \in \mathbb{R}^n$ generally requires $\mathcal{O}(mn)$ operations, which can be superlinear in n . For even modest values m (e.g., when $m = n^\beta$ for $\beta \in (0, 1]$ the complexity is $\mathcal{O}(n^{1+\beta})$). Here, we seek computationally-efficient approximations of \mathbf{A} implementable via convolution and downsampling (and, perhaps, modest post-processing).

We first consider approximating \mathbf{A} as $\mathbf{A} \approx \mathbf{S}\mathbf{C}$, where \mathbf{C} is an $n \times n$ circulant matrix and \mathbf{S} is a matrix with m rows that are a (permuted) subset of \mathbf{I}_n ($n \times n$ identity). These *partial circulant* approximations enjoy an implementation complexity of $\mathcal{O}(n \log n)$, owing to fast Fourier transform methods. Despite these potential implementation complexity benefits, our first contribution establishes that most LDR matrices, in fact, *cannot* be well-approximated by partial circulant matrices.

We then propose a generalization that uses approximations of the form $\mathbf{A} \approx \mathbf{P}\mathbf{S}\mathbf{C}$, where \mathbf{C} and \mathbf{S} are as above, except that \mathbf{S} has some $m' \geq m$ rows, and \mathbf{P} is an $m \times m'$ “post-processing” matrix. Operating with such matrices requires $\mathcal{O}(mm' + n \log n)$ operations in general, and can be as low as $\mathcal{O}(n \log n)$, e.g., when $m' = \mathcal{O}(n^{1/2})$. Within this framework, we also exploit the fact that signals of interest often reside on a *restricted input domain* (e.g., they lie on a union of subspaces, manifold, etc.), so we may restrict our approximation to mimicking the action of \mathbf{A} on these inputs. We describe a *data-driven* approach to learning the factors of the approximating matrix in these settings, and provide empirical evidence to demonstrate the viability of this approximation approach.

B. Connections to Existing Work

Circulant approximations to square matrices are classical in linear algebra; for example, circulant preconditioners for linear systems were examined in [33]–[35]. In a line of work motivated by “optical information processing,” several efforts have examined fundamental aspects of approximating square matrices by products of circulant and diagonal matrices [36]–[40]. Here, our focus is on LDR matrices (not square matrices), so results from these works are not directly applicable here.

Random partial circulant matrices have been studied recently in compressed sensing tasks [41]–[45], and random partial circulant matrices with diagonal pre-processors have been proposed as computationally efficient methods for Johnson-Lindenstrauss (JL) embeddings in [46], [47]. The work [48] established the viability of using random partial circulant matrices for embedding manifold-structured data. In contrast to these works, here our aim is to approximate the action of *given* LDR matrix, not necessarily to perform JL embeddings.

C. Outline

Following the introduction of a few preliminaries (below), we establish in Sect. II a fundamental approximation result regarding partial circulant approximations to general LDR matrices. In Sect. III we propose a generalized approach to the partial circulant approximation problem, and provide empirical evidence to demonstrate its efficacy. A brief discussion is provided in Sect. IV; proofs appear in the Appendix.

D. Preliminaries

We introduce some preliminary concepts and notation that will be used throughout the sequel. First, we define

$$\mathbf{R} = \begin{bmatrix} \mathbf{0}_{(n-1) \times 1} & \mathbf{I}_{n-1} \\ 1 & \mathbf{0}_{1 \times (n-1)} \end{bmatrix},$$

to be the “right rotation” matrix, in that post-multiplication of a row vector by \mathbf{R} implements a circular shift to the right by one position. Analogously, post-multiplying a row vector by \mathbf{R}^T implements a circular shift to the left by one position. To simplify notation we write $\mathbf{L} = \mathbf{R}^T$; note that $\mathbf{L}\mathbf{R} = \mathbf{I}_n$.

We represent an $n \times n$ (real) circulant matrix by

$$\mathbf{C} = \begin{bmatrix} c_1 & c_2 & \cdots & c_n \\ c_n & c_1 & \cdots & c_{n-1} \\ & & \ddots & \\ c_2 & c_3 & \cdots & c_1 \end{bmatrix} = \begin{bmatrix} \mathbf{c}^T \\ \mathbf{c}^T \mathbf{R} \\ \vdots \\ \mathbf{c}^T \mathbf{R}^{n-1} \end{bmatrix}, \quad (1)$$

where $\mathbf{c} = [c_1 \ \cdots \ c_n]^T \in \mathbb{R}^n$. We let \mathcal{C}_n denote the set of all (real) n -dimensional circulant matrices of the form (1). For $m < n$, the set of all $m \times n$ real *partial* circulant matrices is

$$\mathcal{PC}_{m,n} = \{\mathbf{S}\mathbf{C} \in \mathbb{R}^{m \times n} \mid \mathbf{S} \in \mathcal{S}_m, \mathbf{C} \in \mathcal{C}_n\},$$

where \mathcal{S}_m is the set of $m \times n$ row sampling matrices whose rows comprise m different canonical basis vectors of \mathbb{R}^n .

For a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, we denote its m individual rows by $\mathbf{A}_{i,:}$ for $i = 1, 2, \dots, m$ and its n columns by $\mathbf{A}_{:,j} \in \mathbb{R}^m$ for $j = 1, \dots, n$. We write the squared Frobenius norm of \mathbf{A} as $\|\mathbf{A}\|_F^2 = \sum_{i,j} |A_{i,j}|^2$, and its 1, 2 norm by $\|\mathbf{A}\|_{1,2} = \sum_{j=1}^n \|\mathbf{A}_{:,j}\|_2$, where $\|\mathbf{A}_{:,j}\|_2$ denotes the Euclidean norm of $\mathbf{A}_{:,j}$.

II. A FUNDAMENTAL (NEGATIVE) APPROXIMATION RESULT FOR PARTIAL CIRCULANT MATRICES

As above, we let $\mathbf{A} \in \mathbb{R}^{m \times n}$ denote the LDR matrix we aim to approximate using a partial circulant matrix of the same dimensions. Here, we consider a tractable (and somewhat natural) choice of the approximation error metric, and seek to find the matrix $\mathbf{Z} \in \mathcal{PC}_{m,n}$ closest to \mathbf{A} in the Frobenius sense. In this setting, the minimum approximation error is given by

$$\mathcal{E}_{\mathcal{PC}_{m,n}}(\mathbf{A}) = \min_{\mathbf{Z} \in \mathcal{PC}_{m,n}} \|\mathbf{A} - \mathbf{Z}\|_F^2. \quad (2)$$

Finding the minimizer of (2) is a non-convex optimization problem, owing to the product nature of elements of $\mathcal{PC}_{m,n}$. Nevertheless, we obtain a precise characterization of the minimum achievable approximation error for a given \mathbf{A} via the following lemma (whose proof appears in the appendix).

Lemma 2.1: For $\mathbf{A} \in \mathbb{R}^{m \times n}$, we have

$$\mathcal{E}_{\mathcal{PC}_{m,n}}(\mathbf{A}) = \|\mathbf{A}\|_F^2 - R^2(\mathbf{A}), \quad (3)$$

where

$$R(\mathbf{A}) = \max_{\mathbf{f} \in \mathcal{F}} \frac{\|\sum_{i=1}^m \mathbf{A}_{i,:} \mathbf{L}^{\mathbf{f}_i}\|_2}{\sqrt{m}}, \quad (4)$$

and $\mathcal{F} = \left\{ \mathbf{f} = [f_1 \dots f_m] \in \{0, \dots, n-1\}^m \mid f_i \neq f_j \ \forall i \neq j \right\}$.

We call the term $R(\mathbf{A})$ in (4) the *Rubik's Score* of the matrix \mathbf{A} , inspired by the fact that $R(\mathbf{A})$ is maximized when the circular shifts of the rows $\{\mathbf{A}_{i,:}\}$ are “maximally aligned” (indeed, the numerator of (4) is a sum of rotated rows of \mathbf{A}).

Evidently, $0 \leq R(\mathbf{A}) \leq \|\mathbf{A}\|_F$, with $R(\mathbf{A}) = \|\mathbf{A}\|_F$ for all $\mathbf{A} \in \mathcal{PC}_{m,n}$, but beyond these “boundary” cases, Lemma 2.1 yields limited interpretational insight into the achievable error for any particular \mathbf{A} . We gain additional insight here using a probabilistic technique – instead of quantifying the approximation error for a fixed \mathbf{A} , we consider instead matrices \mathbf{A} whose row spaces are distributed uniformly at random on $\text{Gr}(m, n)$, the Grassmannian manifold of m -dimensional linear subspaces of \mathbb{R}^n . Using this model, we can quantify the *proportion* of matrices so drawn whose optimal partial circulant approximation error is at most a fixed fraction (say δ) of their squared Frobenius norm.

To this end, we exploit the fact that matrices whose row-spaces are uniformly distributed on $\text{Gr}(m, n)$ may be modeled as matrices having iid zero-mean Gaussian elements. With this, we establish the following theorem (proved in the appendix).

Theorem 2.1: For $2 \leq m \leq n$, let $\mathbf{A} \in \mathbb{R}^{m \times n}$ have iid $\mathcal{N}(0, 1)$ entries. Then for $\delta \in [0, 0.125)$, and n is sufficiently large, there exists a positive constant $c(\delta)$ such that

$$\Pr(\mathcal{E}_{\mathcal{PC}_{m,n}}(\mathbf{A}) \leq \delta \|\mathbf{A}\|_F^2) = \mathcal{O}(e^{-c(\delta) \cdot mn}).$$

Algorithm 1 “Data-Driven” Partial Circulant Approximation**Inputs:** LDR matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, parameters $\lambda, \mu, \epsilon > 0$,Matrix of “representative” data $\mathbf{X} \in \mathbb{R}^{n \times p}$,**Initialize:** $\mathbf{M}^{(0)} = \mathbf{U}\Sigma$ (from the SVD $\mathbf{A}\mathbf{X} = \mathbf{U}\Sigma\mathbf{V}^T$)

$$\text{obj}^{(0)} = \|\mathbf{A}\mathbf{X}\|_F^2$$

repeat

$$\mathbf{C}^{(t)} = \arg \min_{\mathbf{C} \in \mathcal{C}_n} \|\mathbf{A}\mathbf{X} - \mathbf{M}^{(t-1)}\mathbf{C}\mathbf{X}\|_F^2 + \mu\|\mathbf{C}\|_F^2$$

$$\mathbf{M}^{(t)} = \arg \min_{\mathbf{M} \in \mathbb{R}^{m \times n}} \|\mathbf{A}\mathbf{X} - \mathbf{M}\mathbf{C}^{(t)}\mathbf{X}\|_F^2 + \lambda\|\mathbf{M}\|_{2,1}$$

$$\text{obj}^{(t)} = \|\mathbf{A}\mathbf{X} - \mathbf{M}^{(t)}\mathbf{C}^{(t)}\mathbf{X}\|_F^2 + \mu\|\mathbf{C}^{(t)}\|_F^2 + \lambda\|\mathbf{M}^{(t)}\|_{2,1}$$

until $\text{obj}^{(t)} - \text{obj}^{(t-1)} \leq \epsilon \cdot \text{obj}^{(t-1)}$ **Output:** $\mathbf{M}^* = \mathbf{M}^{(t)}$, $\mathbf{C}^* = \mathbf{C}^{(t)}$

Simply put, the content of Theorem 2.1 is that the proportion of large matrices (with uniformly distributed row spaces) that can be approximated to high accuracy by partial circulant matrices is exponentially small in the product of the matrix dimensions. In the next section we propose a more general framework designed to facilitate accurate partial circulant approximations in a number of practical applications.

III. A GENERALIZED “DATA-DRIVEN” PARTIAL CIRCULANT APPROXIMATION APPROACH

That the conclusion of Theorem 2.1 be pessimistic is, perhaps, intuitive, since $\mathbf{A} \in \mathbb{R}^{m \times n}$ has $\mathcal{O}(mn)$ “degrees of freedom,” while any $\mathbf{C} \in \mathcal{PC}_{m,n}$ has only $\mathcal{O}(n+m)$. These structural limitations are, of course, fundamental. We gain additional leverage here by exploiting a more general approximation model, and exploiting additional latent “signal space characteristics” that arise in many practical applications.

Formally, we consider approximations of \mathbf{A} of the form $\mathbf{A} \approx \mathbf{PSC}$ where $\mathbf{C} \in \mathcal{C}_n$, \mathbf{S} is an $m' \times n$ matrix comprising a permuted subset of rows of identity, and $\mathbf{P} \in \mathbb{R}^{m \times m'}$ is a “post-processing” matrix. This extension allows approximations of \mathbf{A} whose rows are generally *linear combinations* of the rows of a partial circulant matrix rather than being rows of a partial circulant matrix themselves, facilitating accurate approximation of any matrix whose row space is *spanned by* vectors related by circular shifts.

Further, we exploit the fact that in many applications where LDR methods are deployed (e.g. PCA, LDA, Compressive Sensing, etc.), the high-dimensional data to be processed are not arbitrary, but instead lie in some *restricted input domain* (e.g., they lie in a low-dimensional subspace, a union of low-dimensional subspaces, distinct clusters, etc.). In these cases, our approximation task reduces to a simpler task of mimicking the action of \mathbf{A} on these restricted inputs.

Suppose $\mathbf{X} \in \mathbb{R}^{n \times p}$ is a matrix whose columns are “representative” of the restricted input domain for the problem of interest. Then, minimizing an objective of the form $\|\mathbf{A}\mathbf{X} - \mathbf{PSCX}\|_F^2$ essentially ensures the approximation be accurate in an (empirical) *average* sense. We propose an algorithmic approach based on alternating minimization and regularized matrix factorization for identifying the factors comprising this more general approximation; the method

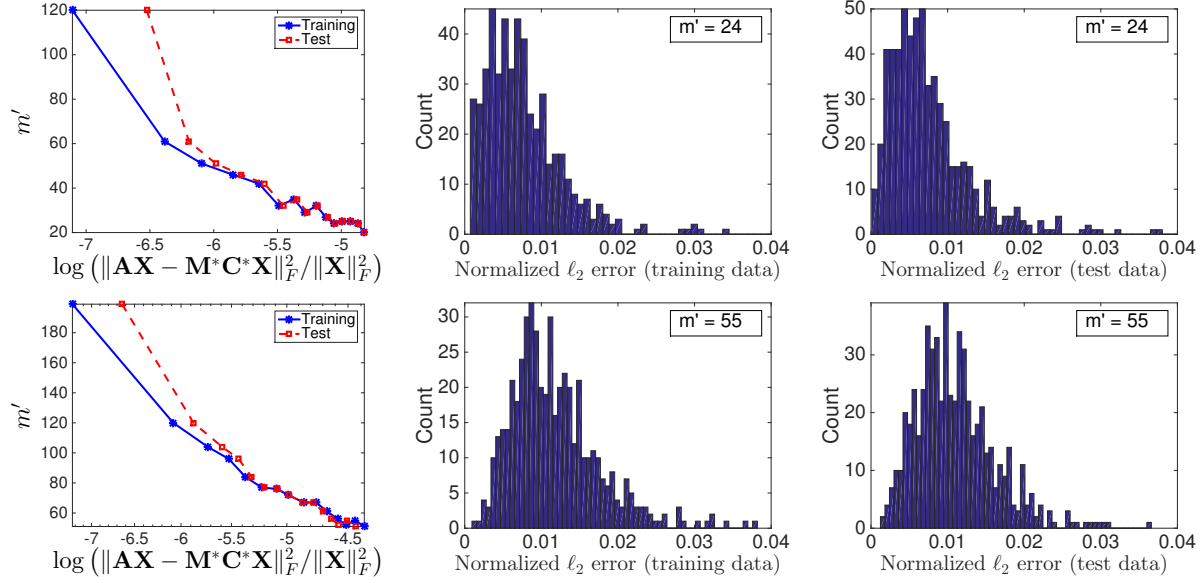


Fig. 1. Data-driven approximation of matrices comprised of top principal components of a training set of ‘2’ digits from the USPS dataset. The panels are (from left to right) average approximation error vs. m' , histogram of normalized training errors, and histogram of normalized test data errors. The top and bottom rows correspond to matrices \mathbf{A} whose rows are the top 5 and top 10 principal component direction vectors, respectively.

is summarized as Algorithm 1. Note that in our approach, we effectively combine the action of the sampling and post processing matrices; we let $\mathbf{PS} \triangleq \mathbf{M}$, and seek *column sparsity* in \mathbf{M} using the $\|\mathbf{M}\|_{2,1}$ regularization term, which penalizes the number of nonzero columns of \mathbf{M} (the resulting number m' of nonzero columns depends on the regularization parameter $\lambda > 0$). We also include a (simple) regularization term for $\mathbf{C} \in \mathbb{R}^{n \times n}$ to fix scaling ambiguities; leaving \mathbf{C} unconstrained would allow elements of \mathbf{M} to become arbitrarily small, and would circumvent the effect of the column-sparsity regularization. Both of these sub-problems are convex, and can be solved using off-the-shelf software (e.g., SLEP [49]).

To evaluate the performance of this approach, we chose \mathbf{X} to be a matrix of training data whose columns are 500 vectorized 16×16 images of the digit ‘2’ from the USPS handwritten digit dataset (available at <http://www.cs.nyu.edu/~roweis/data.html>), and held out the remaining images of the digit ‘2’ to form a “test” data set. We select the rows of the matrix \mathbf{A} to be a subset of the top principal component vectors of the training set \mathbf{X} , and we consider two settings corresponding to \mathbf{A} having 5 rows and 10 rows. Then, we applied the procedure of Alg. 1, with $\mu = 0.1$ and varying values of λ to obtain different factorizations \mathbf{M}^* and \mathbf{C}^* . For each, we quantified the normalized error on the training set (on a logarithmic scale) vs. the column sparsity of \mathbf{M}^* . For a convenient choice of λ , we also computed histograms of the normalized approximation errors ($\|\mathbf{Ax} - \mathbf{MCx}\|_2^2 / \|\mathbf{x}\|_2^2$) for each point in the testing and training data sets. The results are provided in Fig. 1, and illustrate that accurate approximations can indeed be obtained for reasonable values of m' for each case, and that the approximations generalize well to the test data, in that the approximation error histogram is concentrated around small values for each ($\leq 2\%$ error

for most points).

IV. DISCUSSION AND CONCLUSIONS

In this work we established a fundamental result regarding partial circulant approximations of general matrices, and proposed and demonstrated a general approach for approximating the action of general linear dimensionality reduction (LDR) operators over restricted input domains. The convolutional nature of the approximations we consider here, in addition to their computational expediency, makes them viable for implementation using LTI systems, or for specialized applications (e.g., RADAR) where convolutional models arise naturally.

A more general investigation motivated by computational complexity considerations would also include other specialized matrix structures with low implementation complexities (e.g., sparse matrices or “fast” or sparse Johnson-Lindenstrauss embeddings [50], [51]). A thorough, unifying analysis of such “low-complexity” approximations is the topic of our ongoing work, and will be reported in a future publication.

APPENDIX

A. Proof of Lemma 2.1

We first write $\mathcal{E}_{\mathcal{P}C_{m,n}}(\mathbf{A})$ in (2) equivalently as

$$\mathcal{E}_{\mathcal{P}C_{m,n}}(\mathbf{A}) = \min_{\mathbf{S} \in \mathcal{S}_m} \min_{\mathbf{C} \in \mathcal{C}_n} \|\mathbf{A} - \mathbf{S}\mathbf{C}\|_F^2. \quad (5)$$

Now, a key insight is that each choice of $\mathbf{S} \in \mathcal{S}_m$ corresponds to a vector \mathbf{f} from the set \mathcal{F} defined in Lemma 2.1, as including the i -th row of \mathbf{I}_n in \mathbf{S} corresponds to selecting the i -th row of \mathbf{C} , which is given from (1) by $\mathbf{c}^T \mathbf{R}^{i-1}$. Using this we may parameterize the choice of \mathbf{S} in terms of \mathbf{f} , and rewrite the objective function in (5) as

$$\begin{aligned} \|\mathbf{A} - \mathbf{S}\mathbf{C}\|_F^2 &= \sum_{i=1}^m \|\mathbf{A}_{i,:} - \mathbf{c}^T \mathbf{R}^{f_i}\|_2^2 \\ &= \sum_{i=1}^m \|\mathbf{A}_{i,:}\|_2^2 + \mathbf{c}^T \mathbf{R}^{f_i} \mathbf{L}^{f_i} \mathbf{c} - 2\mathbf{A}_{i,:} \mathbf{L}^{f_i} \mathbf{c}. \end{aligned}$$

Thus, since $\mathbf{R}^{f_i} \mathbf{L}^{f_i} = \mathbf{I}_n$, we have

$$\mathcal{E}_{\mathcal{P}C_{m,n}}(\mathbf{A}) = \min_{\mathbf{f} \in \mathcal{F}} \min_{\mathbf{c} \in \mathbb{R}^n} \|\mathbf{A}\|_F^2 + m\mathbf{c}^T \mathbf{c} - 2 \left(\sum_{i=1}^m \mathbf{A}_{i,:} \mathbf{L}^{f_i} \right) \mathbf{c}. \quad (6)$$

We first minimize the objective with respect to \mathbf{c} , keeping \mathbf{f} fixed to any arbitrary value in \mathcal{F} . This is an unconstrained strictly convex quadratic problem whose minima can be obtained by equating the gradient (with respect to \mathbf{c}) to zero. We identify the optimal \mathbf{c} to be $\mathbf{c}^* = \frac{1}{m} \sum_{i=1}^m \mathbf{R}^{f_i} (\mathbf{A}_{i,:})^T$. Substituting \mathbf{c}^* into (6), and simplifying, we obtain

$$\mathcal{E}_{\mathcal{P}C_{m,n}}(\mathbf{A}) = \min_{\mathbf{f} \in \mathcal{F}} \|\mathbf{A}\|_F^2 - \frac{1}{m} \left\| \sum_{i=1}^m \mathbf{A}_{i,:} \mathbf{L}^{f_i} \right\|_2^2,$$

which gives (3), using the definition (4) of the *Rubik's Score*.

B. Proof of Theorem 2.1

Note that $|\mathcal{F}| = n!/(n-m)!$ (since it is just the number of ways of choosing m elements out of n without replacement), and we have (trivially) that $n!/(n-m)! < n^m$, so

$$\begin{aligned} \Pr(\mathcal{E}_{\mathcal{P}_{C,m,n}}(\mathbf{A}) \leq \delta \|\mathbf{A}\|_F^2) &= \Pr\left((1-\delta)\|\mathbf{A}\|_F^2 \leq \sup_{f \in \mathcal{F}} \frac{\|\sum_{i=1}^m \mathbf{A}_{i,:} \mathbf{L}^{f_i}\|_2^2}{m}\right) \\ &= \Pr\left(\bigcup_{f \in \mathcal{F}} \left[(1-\delta)\|\mathbf{A}\|_F^2 \leq \frac{\|\sum_{i=1}^m \mathbf{A}_{i,:} \mathbf{L}^{f_i}\|_2^2}{m}\right]\right) \\ &\leq n^m \Pr\left((1-\delta)\|\mathbf{A}\|_F^2 \leq \frac{\|\sum_{i=1}^m \mathbf{A}_{i,:} \mathbf{L}^{f_i}\|_2^2}{m}\right), \end{aligned}$$

where the last step follows from union bounding.

Now, we can simplify further by introducing the shorthand notation $\mathbf{a} = [\mathbf{A}_{1,:} \ \mathbf{A}_{2,:} \ \dots \ \mathbf{A}_{m,:}]^T \in \mathbb{R}^{mn}$ and $\tilde{\mathbf{R}}_{\mathbf{f}} = [\mathbf{R}^{f_1} \ \mathbf{R}^{f_2} \ \dots \ \mathbf{R}^{f_m}] \in \mathbb{R}^{n \times mn}$, so that $\|\mathbf{A}\|_F^2 = \mathbf{a}^T \mathbf{a}$ and $\|\sum_{i=1}^m \mathbf{A}_{i,:} \mathbf{L}^{f_i}\|_2^2/m = \mathbf{a}^T (\tilde{\mathbf{R}}_{\mathbf{f}}^T \tilde{\mathbf{R}}_{\mathbf{f}}/m) \mathbf{a}$. It is easy to check that $(\tilde{\mathbf{R}}_{\mathbf{f}}^T \tilde{\mathbf{R}}_{\mathbf{f}}/m) \tilde{\mathbf{R}}_{\mathbf{f}}^T = \tilde{\mathbf{R}}_{\mathbf{f}}^T$, which implies that the columns of $\tilde{\mathbf{R}}_{\mathbf{f}}^T$ are eigenvectors of the matrix corresponding to eigenvalue 1. Further, since $(\tilde{\mathbf{R}}_{\mathbf{f}}^T \tilde{\mathbf{R}}_{\mathbf{f}}/m)$ is symmetric, it admits an eigendecomposition $(\tilde{\mathbf{R}}_{\mathbf{f}}^T \tilde{\mathbf{R}}_{\mathbf{f}}/m) = \mathbf{U}_{\mathbf{f}} \Sigma_{\mathbf{f}} \mathbf{U}_{\mathbf{f}}^T$ with $\mathbf{U}_{\mathbf{f}}$ orthonormal and $\Sigma_{\mathbf{f}}$ diagonal, and since it is rank n , $\Sigma_{\mathbf{f}}$ has exactly n entries being 1 (and the rest 0).

Incorporating this insight into the probabilistic analysis above, we obtain that

$$\begin{aligned} \Pr\left((1-\delta)\|\mathbf{A}\|_F^2 \leq \frac{\|\sum_{i=1}^m \mathbf{A}_{i,:} \mathbf{L}^{f_i}\|_2^2}{m}\right) &= \Pr(\mathbf{a}^T \mathbf{U}_{\mathbf{f}} ((1-\delta)\mathbf{I}_{mn} - \Sigma_{\mathbf{f}}) \mathbf{U}_{\mathbf{f}}^T \mathbf{a} \leq 0) \\ &= \Pr(\tilde{\mathbf{a}}^T ((1-\delta)\mathbf{I}_{mn} - \Sigma_{\mathbf{f}}) \tilde{\mathbf{a}} \leq 0), \end{aligned}$$

where the components of $\tilde{\mathbf{a}} = \mathbf{U}_{\mathbf{f}} \mathbf{a}$ are iid $\mathcal{N}(0, 1)$ due to the unitary invariance of the Gaussian distribution. Thus, with a slight overloading of notation, we may write that

$$\begin{aligned} \Pr(\mathcal{E}_{\mathcal{P}_{C,m,n}}(\mathbf{A}) \leq \delta \|\mathbf{A}\|_F^2) &\leq n^m \Pr\left(\sum_{i=2}^m \|\mathbf{A}_{i,:}\|_2^2 \leq \delta \sum_{i=1}^m \|\mathbf{A}_{i,:}\|_2^2\right). \end{aligned} \quad (7)$$

At this point, we note that vectorizing (row-wise) a random matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ having iid zero-mean Gaussian elements yields an mn -dimensional vector whose direction is selected uniformly at random from $\text{Gr}(1, mn)$. Thus, $\sum_{i=1}^m \|\mathbf{A}_{i,:}\|_2^2$ quantifies the length of the vector, while $\sum_{i=2}^m \|\mathbf{A}_{i,:}\|_2^2$ describes the energy retained after projecting the vector onto the fixed $n(m-1)$ -dimensional subspace spanned by the last $n(m-1)$ coordinates. It follows that the probability on the right-hand side of (7) may be interpreted in terms of the energy retained after projecting a *fixed* mn -dimensional unit-normed vector onto a subspace selected uniformly at random from $\text{Gr}(n, mn)$. To quantify this, we use the following.

Lemma A.1 (Adapted from Thm. 2.14 of [52]): Let \mathbf{v} be a fixed unit-length vector in \mathbb{R}^d , W a randomly oriented k -dimensional subspace, and \mathbf{w} the projection of \mathbf{v} onto W . For $0 \leq \varepsilon \leq 1$,

$$\Pr \left(\|\mathbf{w}\|_2 \leq (1 - \varepsilon)\sqrt{k/d} \right) \leq 3e^{-k\varepsilon^2/64}.$$

Using this result (with $k = n(m-1)$ and $d = mn$) we obtain that for $\delta < 1 - 1/m$,

$$\begin{aligned} \Pr(\mathcal{E}_{\mathcal{PC}_{m,n}}(\mathbf{A}) \leq \delta \|\mathbf{A}\|_F^2) &\leq n^m \Pr \left(\frac{\sum_{i=2}^m \|\mathbf{A}_{i,:}\|_2^2}{\sum_{i=1}^m \|\mathbf{A}_{i,:}\|_2^2} \leq \delta \right) \\ &\leq 3n^m e^{-\frac{n(m-1)}{64} \left(1 - \sqrt{\frac{\delta}{1-\frac{1}{m}}}\right)^2} \\ &\leq 3 e^{m \log n - \frac{n(m-1)}{64} \left(1 - 2\sqrt{\frac{\delta}{1-\frac{1}{m}}}\right)}. \end{aligned}$$

Now, it is straightforward to verify that for $\delta < 1/8 = 0.125$ and n sufficiently large, so that $n/\log n > 128/(1 - 2\sqrt{2\delta})$, there exists a positive constant

$$c(\delta) < \frac{1}{128} \left(1 - 2\sqrt{2\delta}\right) - \frac{\log n}{n}$$

for which

$$m \log n - \frac{n(m-1)}{64} \left(1 - 2\sqrt{\frac{\delta}{1-\frac{1}{m}}}\right) \leq -c(\delta)mn.$$

In this case, it follows that

$$\Pr(\mathcal{E}_{\mathcal{PC}_{m,n}}(\mathbf{A}) \leq \delta \|\mathbf{A}\|_F^2) \leq 3e^{-c(\delta) \cdot mn}.$$

REFERENCES

- [1] K. Pearson, "On lines and planes of closest fit to systems of points in space," *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 2, no. 11, pp. 559–572, 1901.
- [2] I. Jolliffe, *Principal component analysis*, Wiley Online Library, 2002.
- [3] J. P. Cunningham and Z. Ghahramani, "Unifying linear dimensionality reduction," *arXiv preprint arXiv:1406.0873*, 2014.
- [4] E. J. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal recovery from highly incomplete frequency information," *IEEE Trans. on Inform. Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [5] D. Donoho, "Compressed sensing," *IEEE Trans. on Inform. Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [6] E. J. Candès and T. Tao, "Near optimal signal recovery from random projections: Universal encoding strategies?," *IEEE Trans. on Inform. Theory*, vol. 52, no. 12, pp. 5406–5425, 2006.
- [7] J. Haupt and R. Nowak, "Signal reconstruction from noisy random projections," *IEEE Transactions on Information Theory*, vol. 52, no. 9, pp. 4036–4048, 2006.
- [8] M. Elad, "Optimized projections for compressed sensing," *IEEE Trans. Signal Processing*, vol. 55, no. 12, pp. 5695–5702, 2007.
- [9] M. Rudelson and R. Vershynin, "On sparse reconstruction from Fourier and Gaussian measurements," *Communications on Pure and Applied Mathematics*, vol. 61, no. 8, pp. 1025–1045, 2008.
- [10] A. Ashok, P. K. Baheti, and M. A. Neifeld, "Compressive imaging system design using task-specific information," *Applied Optics*, vol. 47, no. 25, pp. 4457–4471, 2008.
- [11] J. M. Duarte-Carvajalino and G. Sapiro, "Learning to sense sparse signals: Simultaneous sensing matrix and sparsifying dictionary optimization," *IEEE Trans. Image Processing*, vol. 18, no. 7, pp. 1395–1408, 2009.
- [12] R. Calderbank, S. Howard, and S. Jafarpour, "Construction of a large class of deterministic sensing matrices that satisfy a statistical isometry property," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 358–374, 2010.

- [13] J. Haupt, L. Applebaum, and R. Nowak, "On the restricted isometry of deterministically subsampled Fourier matrices," in *IEEE Conf. on Information Sciences and Systems*, 2010, pp. 1–6.
- [14] L. Zelnik-Manor, K. Rosenblum, and Y. C. Eldar, "Sensing matrix optimization for block-sparse decoding," *IEEE Trans. Signal Processing*, vol. 59, no. 9, pp. 4300–4312, 2011.
- [15] A. Ashok and M. A. Neifeld, "Compressive imaging: hybrid measurement basis design," *J. Optical Society of America A*, vol. 28, no. 6, pp. 1041–1050, 2011.
- [16] W. R. Carson, M. Chen, M. R. D. Rodrigues, R. Calderbank, and L. Carin, "Communications-inspired projection design with application to compressive sensing," *SIAM Journal on Imaging Sciences*, vol. 5, no. 4, pp. 1185–1212, 2012.
- [17] S. Jain, A. Soni, and J. Haupt, "Compressive measurement designs for estimating structured signals in structured clutter: A Bayesian experimental design approach," in *Asilomar Conf. on Signals, Systems and Computers*. IEEE, 2013, pp. 163–167.
- [18] S. Ji, Y. Xue, and L. Carin, "Bayesian compressive sensing," *IEEE Trans. on Signal Processing*, vol. 56, no. 6, pp. 2346–2356, 2008.
- [19] J. Haupt, R. Baraniuk, R. Castro, and R. Nowak, "Compressive distilled sensing: Sparse recovery using adaptivity in compressive measurements," in *Proc. Asilomar Conf. on Signals, Systems, and Computers*, 2009, pp. 1551–1555.
- [20] J. Haupt, R. Castro, and R. Nowak, "Adaptive sensing for sparse signal recovery," in *Proc. IEEE DSP Workshop and Workshop on Sig. Processing Education*, 2009, pp. 702–707.
- [21] J. Haupt and R. Nowak, "Adaptive sensing for sparse recovery," in *Compressed Sensing: Theory and applications*, Y. Eldar and G. Kutyniok, Eds. Cambridge University Press, 2011.
- [22] P. Indyk, E. Price, and D. P. Woodruff, "On the power of adaptivity in sparse recovery," in *Proc. IEEE Foundations of Computer Science*, 2011, pp. 285–294.
- [23] A. Ashok, J. L. Huang, and M. A. Neifeld, "Information-optimal adaptive compressive imaging," in *Asilomar Conf. on Signals, Systems and Computers*, 2011, pp. 1255–1259.
- [24] E. Price and D. P. Woodruff, "Lower bounds for adaptive sparse recovery," *arXiv preprint arXiv:1205.3518*, 2012.
- [25] M. Iwen and A. Tewfik, "Adaptive group testing strategies for target detection and localization in noisy environments," *IEEE Trans. on Sig. Processing*, vol. 60, no. 5, pp. 2344–2353, 2012.
- [26] J. Haupt, R. Baraniuk, R. Castro, and R. Nowak, "Sequentially designed compressed sensing," in *Proc. IEEE Statistical Sig. Processing Workshop*, 2012, pp. 401–404.
- [27] M. A. Davenport and E. Arias-Castro, "Compressive binary search," in *Proc. IEEE Intl. Symp. on Inform. Theory*, 2012, pp. 1827–1831.
- [28] E. Arias-Castro, E. J. Candès, and M. A. Davenport, "On the fundamental limits of adaptive sensing," *IEEE Trans. on Inform. Theory*, vol. 59, no. 1, pp. 472–481, 2013.
- [29] J. M. Duarte-Carvajalino, G. Yu, L. Carin, and G. Sapiro, "Task-driven adaptive statistical compressive sensing of gaussian mixture models," *IEEE Trans. Signal Processing*, vol. 61, no. 3, pp. 585–600, 2013.
- [30] A. Soni and J. Haupt, "On the fundamental limits of recovering tree sparse vectors from noisy linear measurements," *IEEE Trans. on Inform. Theory*, vol. 60, no. 1, pp. 133–149, 2014.
- [31] M. Malloy and R. Nowak, "Near-optimal adaptive compressive sensing," *IEEE Trans. on Inform. Theory*, vol. 60, no. 4, pp. 4001–4012, 2014.
- [32] R. M. Castro, "Adaptive sensing performance lower bounds for sparse signal detection and support estimation," *Bernoulli*, vol. 20, no. 4, pp. 2217–2246, 2014.
- [33] G. Strang, "A proposal for Toeplitz matrix calculations," *Studies in Applied Mathematics*, vol. 74, no. 2, pp. 171–176, 1986.
- [34] T. F. Chan, "An optimal circulant preconditioner for Toeplitz systems," *SIAM Journal on Scientific and Statistical Computing*, vol. 9, no. 4, pp. 766–771, 1988.
- [35] E. E. Tyrtyshnikov, "Optimal and superoptimal circulant preconditioners," *SIAM Journal on Matrix Analysis and Applications*, vol. 13, no. 2, pp. 459–473, 1992.
- [36] J. Müller-Quade, H. Aagedal, T. Beth, and M. Schmid, "Algorithmic design of diffractive optical systems for information processing," *Physica D: Nonlinear Phenomena*, vol. 120, no. 1, pp. 196–205, 1998.
- [37] M. Schmid, R. Steinwandt, J. Müller-Quade, M. Rötteler, and T. Beth, "Decomposing a matrix into circulant and diagonal factors," *Linear Algebra and its Applications*, vol. 306, no. 1, pp. 131–143, 2000.
- [38] M. Huhtanen, "How real is your matrix?," *Linear algebra and its applications*, vol. 424, no. 1, pp. 304–319, 2007.

- [39] M. Huhtanen, "Factoring matrices into the product of two matrices," *BIT Numerical Mathematics*, vol. 47, no. 4, pp. 793–808, 2007.
- [40] M. Huhtanen, "Approximating ideal diffractive optical systems," *Journal of Mathematical Analysis and Applications*, vol. 345, no. 1, pp. 53–62, 2008.
- [41] J. Romberg, "Compressive sensing by random convolution," *SIAM Journal on Imaging Sciences*, vol. 2, no. 4, pp. 1098–1128, 2009.
- [42] J. Haupt, W. U. Bajwa, G. Raz, and R. Nowak, "Toeplitz compressed sensing matrices with applications to sparse channel estimation," *IEEE Transactions on Information Theory*, vol. 56, no. 11, pp. 5862–5875, 2010.
- [43] H. Rauhut, "Compressive sensing and structured random matrices," in *Theoretical Foundations and Numerical Methods for Sparse Recovery*, M. Fornasier, Ed., vol. 9, pp. 1–92. deGruyter, 2010.
- [44] H. Rauhut, J. Romberg, and J. A. Tropp, "Restricted isometries for partial random circulant matrices," *Applied and Computational Harmonic Analysis*, vol. 32, no. 2, pp. 242–254, 2012.
- [45] W. Yin, S. Morgan, J. Yang, and Y. Zhang, "Practical compressive sensing with Toeplitz and circulant matrices," in *Visual Communications and Image Processing 2010*. International Society for Optics and Photonics, 2010, p. 77440K.
- [46] F. Krahmer and R. Ward, "New and improved Johnson-Lindenstrauss embeddings via the restricted isometry property," *SIAM Journal on Mathematical Analysis*, vol. 43, no. 3, pp. 1269–1281, 2011.
- [47] H. Zhang and L. Cheng, "New bounds for circulant Johnson-Lindenstrauss embeddings," *arXiv preprint arXiv:1308.6339*, 2013.
- [48] H. L. Yap, M. B. Wakin, and C. J. Rozell, "Stable manifold embeddings with structured random matrices," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 4, pp. 720–730, 2013.
- [49] J. Liu, S. Ji, and J. Ye, *SLEP: Sparse Learning with Efficient Projections*, Arizona State University, 2009.
- [50] N. Ailon and B. Chazelle, "Approximate nearest neighbors and the fast Johnson-Lindenstrauss transform," in *Proc. ACM Symp. on Theory of Computing*, 2006, pp. 557–563.
- [51] A. Dasgupta, R. Kumar, and T. Sarlós, "A sparse Johnson-Lindenstrauss transform," in *Proc. ACM Symp. on Theory of Computing*, 2010, pp. 341–350.
- [52] J. Hopcroft and R. Kannan, *Foundations of Data Science*, April 2014, Available online: <http://www.cs.cornell.edu/jeh/book11April2014.pdf>.